

1 SHEPPARD, MULLIN, RICHTER & HAMPTON LLP
A Limited Liability Partnership
2 Including Professional Corporations
NEIL A.F. POPOVIC, Cal. Bar No. 132403
3 ANNA S. McLEAN, Cal. Bar No. 142233
TENAYA RODEWALD, Cal. Bar No. 248563
4 LIEN H. PAYNE, Cal. Bar No. 291569
JOY O. SIU, Cal. Bar No. 307610
5 Four Embarcadero Center, 17th Floor
San Francisco, California 94111-4109
6 Telephone: 415.434.9100
Facsimile: 415.434.3947
7 Email: npopovic@sheppardmullin.com
amclean@sheppardmullin.com
8 trodewald@sheppardmullin.com
lpayne@sheppardmullin.com
9 jsiu@sheppardmullin.com

10 Attorneys for Defendant,
SEAGATE TECHNOLOGY LLC

12 UNITED STATES DISTRICT COURT

13 NORTHERN DISTRICT OF CALIFORNIA, SAN FRANCISCO DIVISION

15 IN RE SEAGATE TECHNOLOGY LLC
LITIGATION

17 CONSOLIDATED ACTION

Case No. 3:16-cv-00523-JCS

**DECLARATION OF GLEN ALMGREN
IN SUPPORT OF SEAGATE'S
OPPOSITION TO PLAINTIFFS'
MOTION FOR CLASS CERTIFICATION**

Date: March 30, 2018
Time: 9:30 a.m.
Place: Courtroom G
Judge: Hon. Joseph C. Spero

Second Consolidated Amended Complaint
filed: July 11, 2016

24 [UNREDACTED in the PUBLIC RECORD]

DECLARATION OF GLEN ALMGREN

I, Glen Almgren, declare as follows:

1. I have personal knowledge of the facts set forth herein, which are known by me to be true and correct, and if called as a witness, I could and would competently testify thereto.

2. In 1998 I received a B.S. in mechanical engineering from the Colorado School of Mines. I briefly worked as a manufacturing engineer for a company called Howard Industries. I joined Seagate Technology LLC (“Seagate”) in 1999. I started as a product assurance engineer. My role was to test and help conduct failure analysis on new hard disk drives that were in development. In 2003, I became software lead for product assurance which meant I qualified the software that we used for testing hard drives. In 2006 or 2007 I became a reliability program manager (“RPM”). This means I was the reliability engineer or specialist on product core teams. A “core team” is the group responsible for designing and developing a new hard disk drive, and I was the member responsible for all the reliability testing of the drive being developed. I was RPM for Grenada Classic. After I was RPM for Grenada Classic, I became the manager of the RPM group, which means I managed and supervised the work of the other RPMs, including the RPMs for the Grenada BP and Grenada BP2. In 2013 I became the engineering director for Seagate’s reliability group in Longmont, Colorado. In this role, I oversee reliability testing and qualification of Seagate hard drives for production. In other words, my team is responsible for ensuring that hard drives that are being developed meet specifications and reliability goals, and are ready to enter mass production for sale to consumers.

3. I have reviewed the declaration of Andrew Hospodor filed in support of Plaintiffs’ Motion for Class Certification. In it, Mr. Hospodor reviews internal Seagate documents and makes statements or assertions that he claims are based on those documents. However, as explained below, many of Hospodor’s statements are clearly erroneous and/or misinterpret Seagate’s documents.

A. Background on Seagate Drive Testing

4. From 2011 to early 2016, Seagate used the internal codename “Grenada” to refer to a family of drives that included the drives with model number ST3000DM001. The ST3000DM001

1 was the 3 terabyte (TB) version of the Grenada drive. There were three different versions of the
2 Grenada drives: the Grenada Classic, Grenada BP and Grenada BP2.

3 5. Seagate uses reliability or “reliability demonstration” testing (“RDT”) in order to
4 qualify a drive or product for production and sale to customers. During reliability testing, Seagate
5 subjects the drives to testing at high temperatures and workloads, which accelerates the occurrence
6 of potential failures so that Seagate can predict how many failures are likely to occur during normal
7 use in the field. Thus, Seagate subjects the drives to high stress (including maximum workloads and
8 high temperatures) during reliability testing in order to produce in a short amount of time the same
9 failures that would occur during normal use by consumers over a much longer period of time.

10 6. In order to qualify a drive for production, Seagate will test a group of more than
11 1,000 drives for six weeks or longer in RDT. RDT is performed at high work-loads, meaning the
12 drives are forced to continuously read and write large amounts of data. For reliability testing,
13 Seagate transferred 4-5 TB/day of data to each drive, meaning that in our standard 1008 hour (42
14 day, 6 week) test, we transferred 168 to 210 TB total to the drives. Since the Grenada drives were
15 rated for 55 TB/year, this means our testing subjected the drives to the equivalent of 3 to 3.8 years
16 of use. As noted, we also subjected the drives to high temperatures during testing. Test
17 temperatures were 15-20° C hotter than the typical field application. This further accelerated wear
18 on the drives, meaning our testing was covering more than 3-4 years of typical use.

19 7. In addition, Seagate also performed other types of testing to qualify drives for
20 production. These tests include Thermal Cycling, Shock/Vibration, Power Consumption,
21 Load/Unload, Environmental Stress/Storage, Agency Certification, Acoustics.

22 8. As part of reliability testing, we calculate the MTBF (mean time between/before
23 failure¹) and AFR (annualized failure rate) for the drives based on actual test data. The data for
24 calculating MTBF and AFR is the time-to-failure for each drive that failed, and the total time in
25 testing for those drives that didn’t fail during our testing of a specific group of drives. We input this
26 data into industry-standard software (Weibull++ or JMP) for performing reliability analysis

27
28 ¹ In the case of a product like a hard drive that is typically not repaired after it fails, MTBF means the mean time before failure. For products that are repaired, MTBF means the mean time between failure.

1 (Weibull analysis). Using Weibull++ or JMP, Seagate performs a Maximum Likelihood Estimation
2 (“MLE”) to fit a Weibull distribution to Seagate’s actual test data and estimate Beta (the shape
3 parameter) and Eta (the scale parameter).

4 9. Beta and Eta are then used to calculate the projected AFR and MTBF. AFR is
5 calculated from the Weibull parameters and the expected power-on hours (“POH”) per year for the
6 product. For any given test data (from testing a population of drives), the projected AFR will be
7 higher for higher expected POH per year. For example, a test population of drives might have an
8 MTBF of 100,000 and for an expected POH of 2,400 hours/year, the AFR will be approximately
9 2.34%. However, if the same drives have an expected POH of 1,000 hours/year (2-3 hours of use
10 per day), then the AFR will be close to 1.

11 10. In numerous portions of his declaration, Hospodor assumes or claims that Seagate
12 ‘selected,’ ‘assumed’ or utilized a pre-selected Weibull Beta value when calculating AFR for the
13 Grenada drives and products. This is wrong. With regard to these products, Seagate always fit the
14 Weibull distribution and determined the Beta value and AFR from the actual test data for the group
15 of drives at issue using the MLE method as explained above. Seagate never pre-selected or assumed
16 a Beta value. The documents that Hospodor cites state that the AFR is “From all fails Weibull
17 MLE”—clearly indicating that the AFR is obtained from maximum likelihood estimation on the
18 failures during actual testing. (See Figure 1 below, which is an example of one of the documents
19 Hospodor cites.) The documents also report the Weibull Beta value in connection with a specific
20 number of drives tested for a specific number of hours. The Weibull Beta values Hospodor cites
21 change from document to document, and are calculated out to as many as 7 decimal places (*e.g.*,
22 “.4819208,” *see* Ex. 3 [FED_SEAG0026867] at p. 26887). An example of all of these things is
23 shown in Figure 1 below. Finally, I testified in my deposition in this action that we determined AFR
24 and the Weibull parameters (including Beta) based on actual test data. (Ex. 13 [Almgren Depo.] at
25 182:8-183:10, 185:4-186:6.) Hospodor’s belief that Seagate ‘assumed’ or ‘selected’ the Weibull
26 Beta value is wrong.

Grenada BP RDT 3.0			
Grenada BP RDT 3.0 MTBF/FE Table			
AFR (1st year Weibull)	2.942%	From all fails Weibull MLE	Updated: 4/12/12 12:00 AM
MTBF (1st year Weibull)	80365.8		
Minimum AFR:	0.050%	From zero fail Weibull @ 50%	1068 QTY_TESTED
Total Number of Failures	32		
AFR for 1 failure	0.090%	AFR decrease per failure @ 1	2400 POH/Year
			0.4819208 Weibull Beta
			661 Average Test Hours

Figure 1 (Ex. 3 [FED_SEAG0026867] at p. 26887.) (FIGURE 1 SOUGHT TO BE REDACTED)

B. Response to Paragraphs 59-82 of the Hospodor Declaration

11. Seagate first approved the Grenada drive for shipment in a single SBS (Seagate Branded Solutions) external USB product in April 2011. The approval was documented in the “Grenada SBS SAD Final 4/28/11.” (See Ex. 1 [FED_SEAG0026697] (“Grenada SBS SAD Final 4/28/11.”).) The April 2011 SBS SAD approved including the Grenada Classic drive as part of a *single, external USB product* (codename “Rockit”). (Ex. 1 at p. 26699.)

12. The Rockit external USB drive was an external, backup drive product used over a USB connection. The specifications for qualifying the Grenada drive in this product were different than they were when we later released the Grenada drive for use as an internal, desktop hard drive. There was no AFR specification for the external, USB products. Instead, the drive was required to meet an MTBF specification of 100,000 hours, which was appropriate because the drive was going into a low-use product. At the time of the April 2011 SBS SAD, the Grenada drive met Seagate’s reliability specifications for drives used in external, USB products, including an MTBF over 100,000 hours.

13. Seagate calculated the MTBF for the Grenada Classic drives based on actual test data. Exhibit 1 [FED_SEAG0026697] at pages 26704-26705 shows that Seagate tested 1,651 drives for an average of 740 hours each. (This document reports results of combined testing, some of

1 which was shorter than 1008 hours.) Seagate input the actual test and failure data from this test
 2 population into the Weibull++ or JMP software and used Maximum Likelihood Estimation to fit the
 3 Weibull distribution and estimate the Beta and Eta parameters of the Weibull distribution. Based on
 4 these Beta and Eta values, Seagate calculated the MTBF for the drives. Seagate did not “select” or
 5 “assume” the Beta value.

6 14. Furthermore, the drives did not have a projected AFR of over 7% as Hospodor claims
 7 in Paragraphs 65 and 75. The results Hospodor cites (Ex. 1 [FED_SEAG0026697] at pp. 26704-
 8 26705) are results for *combined* Maturity Acceptance Testing (“MAT”) testing 1.2/1.3/BtC/2.0—
 9 not the results of the final testing. Some of these MAT tests occurred earlier in the design and
 10 development process, before the process was complete. Thus, the combined test results do not
 11 represent the state of the drives at the time of qualification. Furthermore, the demonstrated
 12 improved MTBF was over 100,000, meaning that Seagate had demonstrated and implemented fixes
 13 that increased the MTBF to over 100,000, which was what was required to qualify the drive for
 14 shipment in the external, USB (SBS) product.

15 15. In Paragraph 79, Hospodor claims that 23 of the 129 failures listed on pp. 26704-05
 16 of Exhibit 1, were “head instability.” Hospodor suggests that the head instability was due to
 17 contamination or lubrication problems, or some other mechanical problem. He also claims that
 18 “head-related failures, in addition to contamination, was a recurring issue with the Drives.” Head
 19 instability is a specific failure mode that almost always refers to *electrical noise* in the signal from
 20 the read/write head—not mechanical “instability,” lubrication, contaminants or “pile up on the air
 21 bearing surface (“ABS”)” as Hospodor claims.

22 16. Moreover, the numbers of Grenada drives exhibiting head instability in accelerated
 23 reliability testing was small. In the case of the Grenada *Classic* drives at issue in Paragraph 79, the
 24 document Hospodor cites (Ex.1 [FED_SEAG0026697], p. 26704-05), reflects the facts that: (i)
 25 Seagate observed this head instability failure mode in MAT 1.2 and 1.3 testing; (ii) Seagate made
 26 changes to correct the problems; and (iii) Seagate validated that the changes had significantly
 27 reduced the failure mode in MAT 2.0 testing—*before* the drive was approved for shipment to
 28 consumers. (See Ex. 1 [FED_SEAG0026697] at p. 26704.) I explained this in my deposition, and

1 therefore I do not understand why Hospodor did not take into account that three-fourths of the 23
 2 failures he discussed in Paragraph 79 were corrected before the Grenada Classic drive was validated
 3 for SBS. (*See e.g.*, Ex. 13 [Almgren Depo.] at 194:3-9.)

4 17. Similarly, Exhibit 4 [FED_SEAG0026751], page 26757, contains a reliability
 5 summary for Grenada **BP** drives (a different version of the drive). It states that 5 drives out of 1073
 6 drives tested for 1009 hours (42 days) failed due to “head instability.” As explained above,
 7 Seagate’s testing transfers 4-5TB of data per day, so 42 days of testing resulted in approximately
 8 189 TB of data transferred. Since the Grenada drives were rated for 55 TB/year, this is the
 9 equivalent of 3.5 years of use, and it was done at very high temperatures (60° C, which is 140° F).
 10 Thus, a *total* of 5 of 1073 drives (0.4%) failed due to head instability *during the equivalent of 3.5*
 11 *years of use*, which is roughly equivalent to 0.13% per year. Finally, in his Paragraph 116 Hospodor
 12 claims Exhibit 5 [FED_SEAG0057277], p. 57324, shows that “the top ‘issue’ in the above chart [for
 13 Grenada **BP2**] is related to contamination, and the third issue is head-related.” However, that page
 14 shows that 6 of 1058 Grenada **BP2** drives failed because of particle/contamination issues, and the
 15 failure of *1 drive out of 1058* was related to head stability. (One drive out of 1058 drives is
 16 0.095%.) Therefore, while “head instability” was noted as a failure mode for the three versions of
 17 the Grenada drive (Classic, BP and BP2), nearly all of the head instability issues were remediated
 18 *before release* of the Grenada Classic, and the percentage of Grenada BP and BP2 drives with head
 19 instability issues was miniscule. As explained above, “contamination” is also a non-specific term
 20 that can mean any number of *unrelated* phenomena. Because contamination can mean so many
 21 different things, there will always be some contamination-related failures in any group of drives, and
 22 hard drive manufactures are always striving to reduce contamination. Hospodor does not identify
 23 any specific kind of contamination that was a common problem for the Grenada drives.

24 18. On October 18, 2011, Seagate approved the Grenada Classic drive for shipment as an
 25 internal, desktop hard drive. (Ex. 2 [FED_SEAG0026839] (“Grenada SAD Approval 10/18/11.”))
 26 Exhibit 2 [FED_SEAG0026839], page 26844 shows that Seagate tested 1,360 drives for an average
 27 of 953 hours. Seagate then used the Weibull++ or JMP software to perform MLE *on the results of*
 28 *this testing*, to estimate the Beta and Eta parameters and to calculate the projected AFR for the

1 drive. Seagate did not select or assume any value for Beta (or Eta), but rather obtained the values
 2 from MLE on the actual test results on the specific set of 1,360 drives. As Hospodor acknowledges
 3 in Paragraph 80, the “demonstrated reduced AFR” for the drive was 0.95%, meaning that Seagate
 4 had implemented manufacturing changes or improvements demonstrated to provide an AFR of
 5 0.95% prior to releasing the drives for sale to consumers. (See Ex. 2 [FED_SEAG0026839], p.
 6 26844.)

7 19. Hospodor states in Paragraph 80 that “Seagate’s own internal testing of the Drive
 8 revealed an increasing, not decreasing, AFR. AFR typically increases when drives begin to wear
 9 out as they approach *the end of their design life*. *Early wear out* is a hallmark of unreliable products,
 10 and are often due to flaws in the design or components.” These statements do not make sense.
 11 First, Hospodor’s claim that “Seagate’s own internal testing of the Drive revealed an increasing, not
 12 decreasing, AFR” is wrong. Hospodor does not cite to any documents to support this statement, and
 13 it is not supported by anything in Exhibit 2 [FED_SEAG0026839]. Furthermore, it is inconsistent
 14 with my knowledge of the results of reliability testing on the Grenada Classic drive, as explained
 15 throughout this declaration. Second, when Hospodor says that “AFR typically increases when
 16 drives begin to wear out as they approach the *end of their design life*” he appears to be saying that if
 17 a group of drives is used for a long time, when they eventually reach the end of their useful life they
 18 will fail causing the AFR to increase. That would be the opposite of “*early wear out*”—which
 19 would refer to drives failing early in their useful life not at the end of their design life. Third,
 20 Hospodor does not show that his general statements about ‘wear out’ have any relation to the
 21 Grenada Classic drive. They do not.

22 20. Hospodor’s Paragraph 82 is not an accurate or reasonable description of the failure
 23 modes listed in Exhibit 2 [FED_SEAG0026839] to which Hospodor refers. Hospodor discusses
 24 “degraded head,” “particle induced media damage” and “head crashes” as if they were all related
 25 phenomena. Particle induced media damage is a broad category of issues that can have multiple
 26 causes because there are multiple possible sources for particle contamination within a hard drive
 27 (parts wearing and shedding particles over time) or external contamination introduced during
 28 manufacturing. Particle induced media damage does not imply a head crash or a degraded head.

Degraded heads can mean electrical or mechanical problems with the heads, with any number of unrelated causes. In addition, as explained above, unstable heads usually means electrically unstable heads (usually high temperature exacerbated electrical instabilities or noise in the read head). It does not imply a ‘mechanically unstable’ head. Degraded heads and unstable heads do not imply head crashes. Furthermore, if there had been a head crash, the failure description would say head crash. Head crashes were not a prominent failure mode for the Grenada drives, and Hospodor is simply wrong to try to tie together the several disparate types of failures (electrically degraded heads, mechanically degraded heads, particle induced media damage and head crashes) and claim that they are all related. They are not.

C. Response to Paragraphs 83-89 of the Hospodor Declaration

21. Hospodor’s Paragraphs 83 discusses the SBS SAD for the Grenada BP drive on April 18, 2012. (Ex. 3 [FED_SEAG0026867].) The Grenada BP was a new version of the Grenada drive, so it required new validation testing reported in this SBS SAD. However, as in the case of the SBS SAD for the Grenada Classic drives, this SAD approved the Grenada BP for shipment as part of SBS products, which in this case were external, USB products. The AFR reported for the Grenada BP at this time was calculated based on actual test data, specifically, 1068 drives tested for an average of 661 hours at the time of the report. (Ex. 3 [FED_SEAG0026867] at p. 26887.) Using MLE, Seagate estimated the Weibull Beta and Eta parameters from the actual test data, and calculated the reported AFR. However, although Seagate calculated the AFR, Seagate did not advertise an AFR for these products, and the Grenada BP drives were not required to meet an AFR specification to qualify for inclusion in SBS products (external, USB products). As explained above, drives for use in external, USB products were required to meet the 100,000 MTBF specification. Seagate demonstrated a validated MTBF of 118,000 for the Grenada BP drives at the time of this SBS SAD. (See Ex. 3 [FED_SEAG0026867] at p. 26886.)

22. On June 5, 2012, Seagate approved the Grenada BP drive for shipment as an internal, desktop drive (disty/OEM). (See Ex. 4 [FED_SEAG0026751 (“GrenadaBP ECQ Approved Final 6-5-12”)].) By June 5, 2012, Seagate had tested 1068 Grenada BP drives for 1009 hours, and had demonstrated a reduced AFR of 0.98%. (Ex. 4 [FED_SEAG0026751] at p. 26783.) The 2.917%

1 “raw” AFR stated in this document does not include the changes we had demonstrated reduced the
2 AFR to 0.98%.

3 **D. Response to Paragraphs 90-117 and 176 of the Hospodor Declaration**

4 23. Paragraphs 85, 90, 94 and 95, Hospodor claims that he AFR for the Grenada Classic
5 drive “increased” over time “from 2.34% to 2.621% to 3.436%” and that this indicated that Seagate
6 was incorrectly using a Weibull Beta value of less than 1 when Seagate should have been using a
7 Weibull Beta value of greater than 1. Hospodor claims “[t]his β value represents an assumption that
8 the Drive will have a declining failure rate over time.” Hospodor’s assertions reflect a fundamental
9 misunderstanding of failure analysis, the Weibull distribution and parameters, and Seagate’s
10 practices. First, as explained, Seagate did not use a pre-selected, assumed or historical value for
11 Beta. Rather, Seagate used appropriate and standard mathematical techniques (and industry-
12 standard software) to determine the Weibull parameters Beta and Eta based on each new group of
13 test drives. Second, the Weibull distribution describes the distribution of failures in a group of
14 products *as they continue to be used over time—i.e., as the products “age.”* In other words, the
15 Weibull distribution describes the distribution of failures in a population of products from 0 hours *of*
16 *use*, to multiple hours or years *of use*. As an analogy, the Weibull distribution and Beta could be
17 used to describe human mortality rates *as compared to age*—including whether there is a high
18 “infant mortality” and/or whether the rate or probability of death increases with age. To the extent
19 that a Beta of greater than 1 represents an increasing failure rate, it concerns whether there is an
20 increasing rate of failures when a group of products “ages” by being *used* for longer and longer.
21 The numbers Hospodor refers to—2.34% to 2.621% to 3.436%—were predicted *first-year*² failure
22 rates of different groups of drives *manufactured at different times*, but all tested or used for 1008
23 hours or less.³ This sequence of first-year AFRs does not and cannot show whether the failure rate

24
25 ² As Hospodor acknowledges in his footnote 8, the AFR numbers Seagate reported were
26 first-year failure rates because, if Beta is less than 1, then the failure rate in subsequent years will be
27 less than the first-year failure rate. Seagate always determined Beta for the Grenada drives from the
28 actual test data, and it was less than 1, so it is correct that Seagate reported the projected first-year
AFR, which would have been the highest AFR for any year of use of the drives.

³ Hospodor also mixes and matches raw AFR with demonstrated reduced AFR in order to
show his alleged “increasing” failure rate, and the last AFR (3.436%) is from the group of drives
that was tested for the shortest amount of time. Specifically, the first number Hospodor cites, 2.34%

1 of the drives would increase (or decrease) with greater amounts *of age or use*. (And in particular,
 2 the last number cannot show an increasing failure rate with greater amounts of use, because it was
 3 obtained from less testing (less use) than the earlier numbers.) By analogy, comparing the mortality
 4 rate of 40-year-old humans from three different eras would not tell you how the mortality rate
 5 changes as humans age. Thus, the sequence Hospodor cites—2.34% to 2.621% to 3.436%—has
 6 nothing to do with the shape of the Weibull distribution, or the value of Beta.

7 24. In Paragraphs 96-111, Hospodor discusses a report by Andrei Khurshudov, a former
 8 Seagate employee. His report looks at return data for a wide variety of drives, but he did not draw
 9 conclusions about the Grenada drives because there was too little data. In addition, some of
 10 Andrei's statements about Seagate's reliability testing procedures are not accurate. In particular,
 11 Andrei implied that Seagate used a historical or pre-selected, or 'assumed' value of Beta that was
 12 less than 1. Andrei said that Seagate typically used a Beta of 0.55. This is not correct for the vast
 13 majority of Seagate's reliability testing and was not correct in relation to the Grenada drives. As
 14 explained above, we determined Beta for each group of test drives using maximum likelihood
 15 estimation (MLE). Seagate would only use an assumed or historical Beta value if a test population
 16 had fewer than 5 failures—which does not occur for these types of tests in which Seagate tests
 17 hundreds or thousands of drives for hundreds or thousands of hours. However, Andrei was not
 18 involved in reliability testing products for validating or approving them for shipment. Andrei
 19 worked on more theoretical and research projects, not reliability testing of specific products, so
 20 perhaps that is why he didn't know how Seagate actually determined Beta for most products
 21 (including the Grenada drives).

23 was the **demonstrated reduced** AFR from testing 1651 Grenada Classic drives for an average of **740**
 24 **hours** for the April 2011 SBS SAD of Grenada Classic. (Ex. 1 [FED_SEAG0026697] at 26705.)
 25 However, the second number Hospodor cites, 2.621%, was the **raw** AFR from testing 1360 Grenada
 26 Classic drives for an average of **953 hours** for the October 2011 disty and OEM SAD. (Ex. 2
 27 [FED_SEAG0026839] at 26844.) The third number Hospodor cites, 3.436% is the **raw** AFR from
 28 ORT testing an unspecified number of Grenada Classic drives for an average of **535 hours** in June
 2012.. (Ex. 4 [FED_SEAG0026751] at 26785-86.) The actual sequence of demonstrated reduced
 AFRs were 2.34%, 0.95% and 2.35%—which does not show the steady “increase” Hospodor
 claims. (See Ex. 1 at p. 26705, Ex. 2 at p. 26844 and Ex. 4 at p. 26786.) Seagate monitors the
 results of testing throughout the test duration. The documents at issue reported the status of testing
 before the full 1008 hours had been reached, or as in the case of the April 2011 report, combined test
 results some of which were shorter than the full 1008 hours.

25. In addition, Hospodor relies on charts that relate to return rates for products from the “disty” channel. These are drives that are sold through websites or distributors to individual consumers or businesses rather than being sold to large OEMs that pre-install drives inside desktop computers. There are two problems with relying on charts of disty return rates the way Andrei and Hospodor did. First, Seagate observes a lot of consumer abuse of drives in the disty channel. This can be things like mishandling (consumers or shippers drop drives or damage them when installing them into their own computers), or businesses purchase consumer-grade drives and install them in high-vibration, high-workload, 24/7 data centers for which the drives were not designed. As Hospodor acknowledges in his Paragraph 27, the AFR does not include misuse, but the disty channel has a lot of *drive misuse, so it is not a good source of information for discussing AFR*. Second, return rates are not a good approximation of failure rates. A lot of drives are returned for reasons other than failure. In fact, some of Andrei's work showed that *over 75%* of SBS drives that were returned *had nothing wrong with them* and most hadn't ever been used at all.⁴ (See Ex. 15 [FED_SEAG0002320] at 2327 showing that overall, out of 1182 returned FreeAgent GoFlex drives, only 267 had a failure, and the remaining 915 drives (over 75%) were NTF (meaning no trouble found) or DNU (meaning the drive had never even been used).)

26. In addition, Andrei's report was aimed at looking at the relationship between high-workload—the workloads experienced by our enterprise, “near line” (NL), and mission-critical drives—and long-term failure rates. As explained above, our Grenada drives are consumer, desktop drives are rated for 55TB/year workload. Our enterprise products have a much higher workload rating (10x that of Grenada) so they require a more extensive test plan to cover the field conditions they will experience. Our consumer, desktop drives are not intended to be used in those types of applications since most consumers aren't running data centers or similar applications. Andrei did not conclude that there was anything wrong with the way we performed reliability testing on consumer, desktop drives like the Grenada drives.

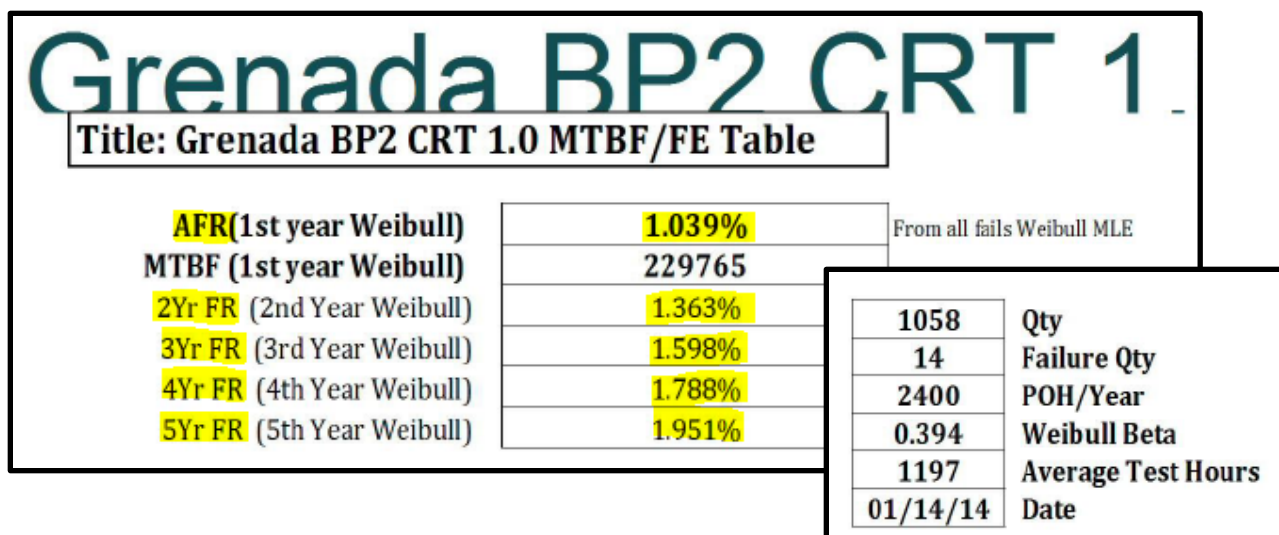
⁴ Although this particular data was for SBS returns, it illustrates that return data is often a very poor predictor of failure rates.

1 27. In fact, Andrei drew no conclusions about Grenada drives—and rightly so. There
 2 simply wasn't enough data. Therefore, Hospodor's discussion of what Andrei did doesn't relate to
 3 whether the Grenada drives had any particular failure rate. It also doesn't show that there was
 4 anything wrong with the way we performed reliability testing on the Grenada drives.

5 28. Finally, Hospodor doesn't ever explain how Seagate's testing is not sufficient, or how
 6 it is not consistent with industry standards. He admits in Paragraph 34 that accelerated life testing is
 7 done for at least 30 days, and Seagate tests its drives for 6 weeks—2 weeks longer. Although he
 8 hints in his Paragraphs 108 and 112 that Seagate should have tested the Grenada drives for longer
 9 than six weeks, he does not explain how longer testing would be appropriate for consumer, desktop
 10 drives or for external, USB drives. As explained above, Seagate's accelerated testing subjects the
 11 drives to the equivalent of *more than 3-4 years* of use and wear. Hospodor doesn't explain how
 12 more extensive testing would be necessary to account for consumer, desktop use or for the even
 13 lower amounts of use for external, USB products. In addition, I recall that we did extended
 14 reliability testing for the Grenada drives when they were qualified in 2012. We subjected drives to a
 15 further 3 or 6 weeks of testing at the same high-workload, high-temperature conditions. We did not
 16 calculate an AFR, but instead used MLE to fit a Weibull distribution to the test data, estimated the
 17 Beta parameter, and confirmed that it continued to remain below 1. I do not recall Seagate's
 18 extended testing ever giving a Beta of greater than 1 for the Grenada drives, and I believe this is
 19 something I would have remembered if it had happened.

20 29. In his Paragraphs 112-115 Hospodor again fundamentally misinterprets a Seagate
 21 document. Hospodor reviews the information shown in Figure 2 below, and claims that it shows the
 22 "AFR" (*annualized* failure rate) for each of years 1 through 5, and specifically that it shows an
 23 "AFR" of 1.039% in the first year and an "AFR" 1.951% in the fifth year. He claims that this shows
 24 an increasing "AFR" and he states that "I cannot reconcile this increasing failure rate with the use of
 25 a Weibull β of 0.394, because a β of <1 reflects a decreasing failure rate." However, in reality,
 26 while the first year is labeled "AFR" and represents a first-year *raw* annualized failure rate of
 27
 28

1 1.039%⁵, the subsequent failure rates are not labeled “AFR” and *do not* represent annualized failure
 2 rates. The subsequent years show total or cumulative failure rates for the corresponding total years
 3 of use. Thus, “2Yr FR” means the projected *total or cumulative failure rate for 2 years of use*.
 4 “3Yr FR” means the total or cumulative failure rate for 3 years of use, etc. Thus, 1.951% for “5Yr
 5 FR” does not mean that 1.951% of drives were projected to fail “in” the 5th year, but instead that
 6 accounting for all the failures over 5 years, Seagate projected that only 1.951% would have failed
 7 (and 98.049% would still be in use) by the end of five years. This further means that the annual
 8 failure rate was decreasing. In the first year, the raw projection (before the demonstrated fixes) was
 9 that 1.039% of drives would fail, but after two years the total failures would only be 1.363%,
 10 meaning that in the 2nd year, the failures were much lower than 1%. In the 3rd year they would be
 11 even lower, etc. In fact, Seagate projected that although 1% of drives would fail in the first year, it
 12 would take another 4 years for the total failures to reach 1.951%. This is fully consistent with the
 13 Beta value of less than 1.



22

23 **Figure 2 (Ex. 5 [FED_SEAG0057277] at p. 57324). (FIGURE 2 SOUGHT TO BE REDACTED)**

24

25 30. Finally, in Paragraph 176, Hospodor again references a number of broad failure
 26 categories (““degraded heads,” “head instability,” and “NMD” (new media defect) caused by


27

28 ⁵ This value, and the others shown in Figure 2, are raw values, before the demonstrated
 changes that reduced the projected first year AFR for this drive to 0.9% as shown in Exhibit 5
 [FED_SEAG0057277] at 57324.

1 contamination”) and tries to claim that they are all related to each other, and implies they are related
2 to head crashes. Again, this is simply not the case. “Head instability” refers to electrical noise in
3 the signal from the head and is not necessarily caused by mechanical degradation of the head or
4 contamination or NMD. “Degraded heads” can refer to mechanical problems *or* electrical problems,
5 and can be caused by many factors other than “abrasion from contacting the disk.” “NMD” may or
6 may not be caused by contamination.

7 I declare under penalty of perjury under the laws of the United States of America that the
8 foregoing is true and correct.

9 Executed on this 4th day of January, 2018, at LONGMONT, Colorado.

10
11 
12 Glen Almgren